

# CO-REGULARIZED DEEP REPRESENTATIONS FOR VIDEO SUMMARIZATION

Olivier Morère<sup>\*,1,2,3</sup>, Hanlin Goh<sup>\*,1,3</sup>, Antoine Veillard<sup>\*,2,3</sup>, Vijay Chandrasekhar<sup>1,3</sup>, Jie Lin<sup>1,3</sup>

I2R<sup>1</sup>, UPMC<sup>2</sup>, IPAL<sup>3</sup>

## ABSTRACT

Compact keyframe-based video summaries are a popular way of generating viewership on video sharing platforms. Yet, creating relevant and compelling summaries for arbitrarily long videos with a small number of keyframes is a challenging task. We propose a comprehensive keyframe-based summarization framework combining deep convolutional neural networks and restricted Boltzmann machines. An original co-regularization scheme is used to discover meaningful subject-scene associations. The resulting multi-modal representations are then used to select highly-relevant keyframes. A comprehensive user study is conducted comparing our proposed method to a variety of schemes, including the summarization currently in use by one of the most popular video sharing websites. The results show that our method consistently outperforms the baseline schemes for any given amount of keyframes both in terms of attractiveness and informativeness. The lead is even more significant for smaller summaries.

**Index Terms**— Video summarization, deep convolutional neural networks, co-regularized restricted Boltzmann machines

## 1. INTRODUCTION

Video sharing websites measure user engagement through click rates and viewership. To make a novel video attractive for the audience, its video link is often presented as a thumbnail of either a single representative frame or a slideshow of several keyframes. In this work, we explore the problem of automatically generating diverse, representative and attractive keyframe-based summaries for videos.

Summarization-based techniques can be broadly divided into three categories: 1) keyframe-based, 2) skimming-based and 3) storyboard-based. In keyframe-based summarization, the video is summarized using a small number of keyframes selected based on some criterion, such as low-level features, like pixel data, motion features, optical flow and frame differences [1, 2, 3], or higher-level information, like objects and



**Fig. 1.** Deep co-regularized keyframe summary. Our method extracts diverse, representative and attractive keyframes.

faces [4, 5]. For this class of algorithms, clustering techniques like k-means are popular: clustering or grouping is performed based on raw RGB pixels, or a combination of low and high level features [6, 7, 8, 9, 10]. The frames closest to the cluster centers are chosen to be part of the summary.

Skimming-based summarization is used to produce longer video summaries. The video is divided into smaller shots using shot boundary detection algorithms and a series of shots are selected to form the summary video. Subshot selection is based on motion activity [11, 12, 13] and other high level features, such as person and landmark descriptors [14].

Finally, in storyboard-based summarization, algorithms take into account relationships between the different subshots [15]. This enables long egocentric videos to be summarized to gain an understanding of the underlying events.

**Contributions.** This work focuses on generating compact keyframe-based summarization, with the main contributions are as follows:

- A comprehensive keyframe-based summarization framework combining deep convolutional neural networks (DCNNs) and restricted Boltzmann machines (RBMs).
- A co-regularization scheme for restricted RBMs able to learn joint high-level subject-scene representations.
- A comprehensive user study comparing our method against various schemes including the algorithm in use by the video sharing website *Dailymotion*.

## 2. CO-REGULARIZED DEEP REPRESENTATIONS

A good keyframe-based summary should consist of easily recognizable subjects in context-setting scenes. To achieve this, we generate frame-level descriptions by exploiting deep

\* O. Morère, H. Goh and A. Veillard contributed equally to this work.

1. Institute for Infocomm Research, A\*STAR, Singapore.

2. Université Pierre et Marie Curie, Paris, France.

3. Image & Pervasive Access Lab, UMI CNRS 2955, Singapore.

convolutional architectures to recognize subjects and scenes. Compact representations are then computed with a novel co-regularization unsupervised learning scheme to exhibit the high-level associations between subjects and scenes. Keyframes are subsequently generated from these compact representations.

## 2.1. Deep Convolutional Neural Networks

Deep convolutional neural networks (DCNNs) have recently been used to obtain astonishing performances in both image classification [16, 17] and image retrieval [18] tasks. For every frame sampled from the video at regular intervals, DCNN descriptors are extracted using the open source Caffe framework [19] along with two pre-trained networks: VGG-ILSVRC-2014-D [20] and Places-CNN [21].

VGG-ILSVRC-2014-D is the best performing single network from the VGG team during the ILSVRC 2014 image classification and localization challenge using the ImageNet [22] dataset. This 138 million parameters network is made of 16 layers: 13 convolutional layers followed by 3 fully-connected layers. It detects 1000 mostly subject-centric categories (e.g. animals, objects, plants, etc.).

Places-CNN is a 60 million parameters network following the AlexNet [16] structure: a total of 8 layers: 5 convolutional layers followed by 3 fully-connected layers. It is trained on the Places 205 dataset, a scene-centric image dataset featuring 205 categories including indoors and outdoors sceneries.


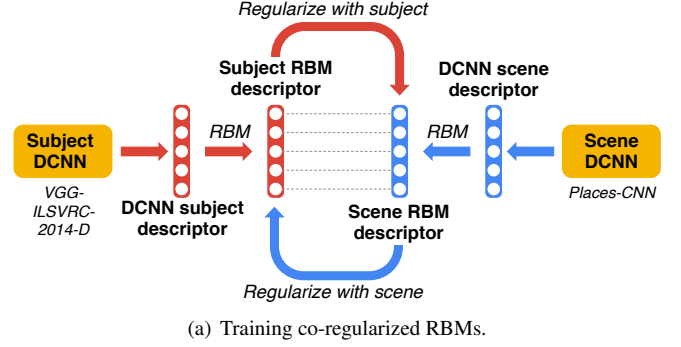
For both DCNNs, descriptors are extracted from the last layer before the softmax operation, having a dimensionality of 1000 and 205 for VGG-ILSVRC-2014-D and Places-CNN, respectively.

## 2.2. Co-Regularized Restricted Boltzmann Machines

To create the video summaries from the DCNN descriptions of subjects  $\mathbf{x}_o$  and scenes  $\mathbf{x}_p$ , we introduce a pair of concurrently trained restricted Boltzmann machines (RBMs) to learn their projections ( $\mathbf{z}_o$  and  $\mathbf{z}_p$ ) to  $K$  units each, where  $K$  is the desired number of keyframes. An RBM is a bipartite network with a projection matrix  $\mathbf{W}$  that maps between its input and output units. RBMs are trained through gradient descent on the approximate maximum likelihood objective, based on network states drawn from Gibbs sampling [23, 24].

In this work, we introduce co-regularization for RBMs. The object RBM is regularized by place representations and in turn regularizes the training of the place RBM (Figure 2(a)). Given randomly sampled minibatches of subject and scene DCNN descriptors  $\{\mathcal{X}_o^i, \mathcal{X}_p^i\}_i$ , we introduce co-regularization cross entropy penalties to the RBM objective functions:

$$\arg \min_{\mathbf{W}_o} - \sum_i \sum_{\mathbf{z}_o^i \in \mathcal{Z}_o^i} \left( \log \mathbb{P}(\mathbf{x}_o^i, \mathbf{z}_o^i) - \lambda_o \sum_k \log \mathbb{P}(\hat{z}_{p,k}^i | \hat{z}_{o,k}^i) \right), \quad (1)$$



**Fig. 2.** A pair of co-regularized RBMs – one representing subjects and another representing scenes – are learned concurrently. (a) During training, an subject unit is regularized by its corresponding scene unit and vice versa. (b) The frame descriptor is a linear combination of the two co-regularized RBM descriptors forming relevant subject-scene associations.

$$\arg \min_{\mathbf{W}_p} - \sum_i \sum_{\mathbf{z}_p^i \in \mathcal{Z}_p^i} \left( \log \mathbb{P}(\mathbf{x}_p^i, \mathbf{z}_p^i) - \lambda_p \sum_k \log \mathbb{P}(\hat{z}_{o,k}^i | \hat{z}_{p,k}^i) \right), \quad (2)$$

where  $\{\mathcal{Z}_o^i, \mathcal{Z}_p^i\}_i$  are the RBM projections of  $\{\mathcal{X}_o^i, \mathcal{X}_p^i\}_i$ ,  $\{\lambda_o, \lambda_p\}$  are the regularization constants, and  $\{\hat{z}_{o,k}^i, \hat{z}_{p,k}^i\}_i$  refer to unit  $k$  in the distribution-sparsified representations of the minibatch [25]. Sparsity across units helps avoid co-adaptation between the units and improves representational diversity across instances of frames. The co-regularization terms serves the purpose of binding a subject and the scene in which it occurs to the same unit position.

The frame descriptor is a linear combination of the two RBM descriptors (Figure 2(b)). The final set of keyframe timings  $t_k, k \in [1..K]$  is the ordered set of  $K$  timings that gives the maximum response for each unit of the frame descriptor:

$$\arg \max_t \alpha z_{o,k}^t + (1 - \alpha) z_{p,k}^t, \quad (3)$$

where  $\alpha \in [0, 1]$  is a balance hyperparameter that causes the summary to be more subject-centric or scene-centric.

This proposed co-regularization method is not specific to subjects or scenes, and is generalizable to other concepts or modalities, such as faces or activities.

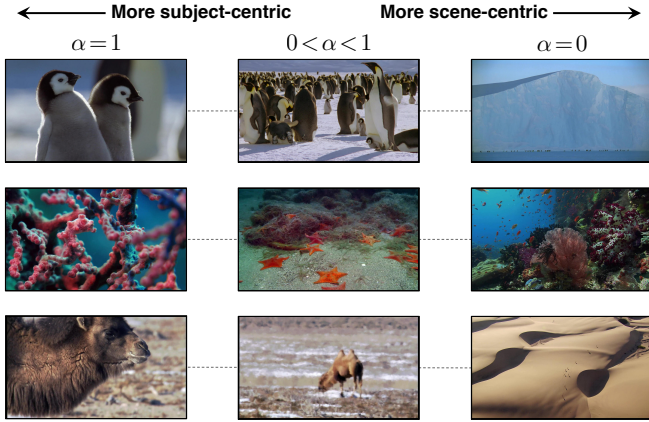
### 3. VIDEO SUMMARIZATION

Using our method, we summarized all 11 episodes from the BBC educational TV series *Planet Earth*<sup>1</sup>. Each episode is approximately 50 minutes long. A sample of our results is shown in Figure 5(a).

#### 3.1. Model Visualization

##### 3.1.1. Balancing Subject- and Scene-Centricity

As shown in Figure 3, bias towards subject or scenes can be adjusted by tuning the  $\alpha$  parameter from Equation 3. This flexibility allows for interesting functionalities such as customising content based on user profiling or explicit queries. The choice of  $\alpha$  value can also be made independently for each unit in order to generate the most visually attractive keyframe, for example based on vibrancy. In practice, setting the default value to  $\alpha = 0.5$  (as used in this empirical study) seems to produce satisfactory results.



**Fig. 3.** Actual keyframes selected by varying  $\alpha$ . Our model can be tuned to select keyframes that are more subject-centric (left), scene-centric (right) or a balance of both (middle).

##### 3.1.2. Visualisation of Co-Regularized RBM Units

Although neural networks tend to be thought of as black boxes, visualization is often useful to decipher what has been learned [26]. To better understand our co-regularized model, we analysed the responses of each unit across the dataset. For this analysis, we trained a single  $K = 12$  model across all 11 episodes. For each of the 24 RBM units, the top 100 frames that most strongly activate each unit were aggregated via a weighted average. The resulting graphical representation of each unit is shown on Figure 4. We observe that the visual appearances of frames corresponding to a subject-scene pair of units are consistently similar. There is also diversity across the units within an RBM.



**Fig. 4.** Visualization of the units for a  $K = 12$  model. The visual representations of subject-scene pairs are well correlated. The categories of the two models are associated in a sensible way and correspond well with the visual representations.

The top 2 categories of each unit identified from the weight matrices are also shown in Figure 4. We notice that the correlation with the visual representation is strong and the subject-scene association is sensibly learned. We can also observe an interesting effect of co-regularization, where associations can be made between subjects (e.g. *polar bear* and *king penguin*) that occur in the same scene (*iceberg*) but never within the same frame.

#### 3.2. User Engagement Study

##### 3.2.1. Evaluation Framework

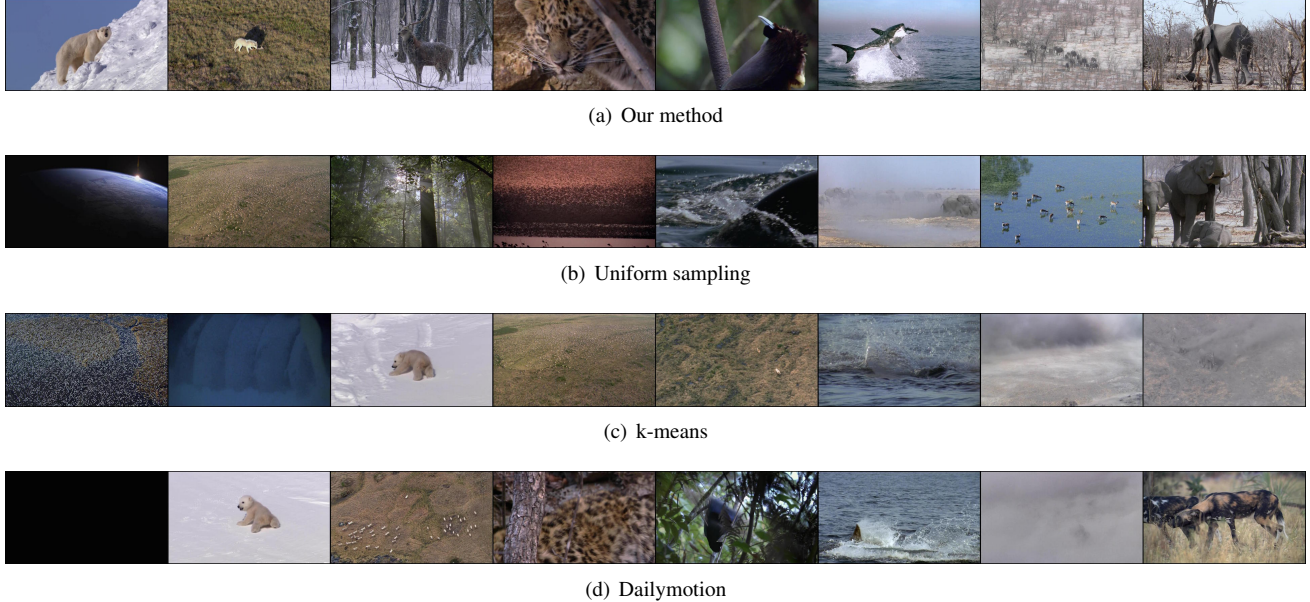
Our method is compared against three other keyframe-based summarisation schemes: naive uniform sampling, k-means clustering and the method currently in use by the video sharing website *Dailymotion*<sup>2</sup>. Each summary is presented as a timeline of keyframes as shown on Figure 5.

Uniform sampling takes  $k$  keyframes with evenly spaced timestamps:  $t_i = \frac{d}{k} (\frac{1}{2} + i)$ ,  $i \in [1..K]$  where  $d$  is the total duration of the video. The k-means clustering scheme uses frames sampled at the same frequency as for our method (1 *fps*) and down-sized to  $32 \times 32$  RGB pixels. Lloyd's algorithm [27] is used to separate the data into  $K$  clusters. 100 runs with different centroid seeds are performed to mitigate the effects of local minima. For each cluster, the frame closest to its centroid is selected as keyframe. *Dailymotion* proposes an 8-keyframes video summary (excluding title frame) which was used as a blackbox scheme to compare our method against. The evaluation videos were uploaded on the website and the proposed summary keyframes were then handpicked from the original footages.

<sup>1</sup><http://www.bbc.co.uk/programmes/b006mywy>

<sup>2</sup><http://www.dailymotion.com/>





**Fig. 5.** Eight keyframes summaries for episode 1 from the TV series *Planet Earth*.

The study was performed by showing pairs of summaries – our method against one of the three baseline schemes – to eight different testers who have not previously seen the videos. For each pair, they are asked to answer the two following questions:

- Q1: Which video would you rather watch? (attractiveness)
- Q2: Which summary was more informative? (informativeness)

Using all the 11 *Planet Earth* episodes, summaries were generated for different amount of keyframes  $K = 4, 6, 8$ , except for *Dailymotion* which imposes  $K = 8$  by default. In total,  $8 \times 11 \times 2 \times 3 \times 8 \times 11 = 616$  answers were collected for each question.

Uniform sampling appears as a natural choice for the wildlife documentaries used during this study given to the slow pace of the action and high visual appeal of the average frame. K-means is expected to be able to capture the diversity of the scenes well whereas it may not perform as well with respect to subjects.

### 3.2.2. Results and Discussion

Table 1 aggregates the answers from the testers. Overall, our method was systematically found more attractive (75% to 97.73% of the time) and more informative (76.14% to 94.32%). Perceived attractiveness and informativeness are strongly correlated. Against *Dailymotion*’s algorithm, our method scores favourably more than three times out of four representing a marked improvement over the scheme currently used by the service.

**Table 1.** How often our method is preferred over each of the three schemes (percentage) for different  $K$ .

$K$	uniform			k-means			daily.
	4	6	8	4	6	8	8
Q1	79.55	82.95	76.14	97.73	82.95	75.00	77.27
Q2	78.41	80.68	81.82	94.32	80.68	76.14	78.41

For varying amounts  $K$  of keyframes, the improvement is rather consistent against uniform sampling whereas against k-means, the improvement is more pronounced when  $K$  is smaller. This is an indication that our overall good method is particularly well-suited for compact summaries.

## 4. CONCLUSIONS

Building upon recent advances in deep learning and image recognition, we proposed a comprehensive keyframe-based summarization framework combining DCNNs and RBMs. Through a comprehensive empirical study, we showed that our method is able to outperform a number of existing schemes. In addition, our novel co-regularization scheme, which discovered meaningful subject-scene associations is generalizable to other concepts and modalities.

Beyond the selection of quality keyframes, our contribution represents a strong step towards the Holy Grail of text-based video summaries by introducing highly interpretable semantic representations.

## ACKNOWLEDGEMENTS

We gratefully acknowledge the support of NVIDIA Corporation with the donation of the GPU used for this research.

## 5. REFERENCES

- [1] T. Liu, H. Zhang, and F. Qi, “A novel video key-frame-extraction algorithm based on perceived motion energy model,” *Circuits and Systems for Video Technology*, vol. 13, no. 10, pp. 1006–1013, 2003.
- [2] W. Wolf, “Key frame selection by motion analysis,” in *Acoustics, Speech, and Signal Processing*, 1996.
- [3] H. Zhang, J. Wu, D. Zhong, and S. W. Smoliar, “An integrated system for content-based video retrieval and browsing,” *Pattern Recognition*, vol. 30, no. 4, pp. 643–658, 1997.
- [4] D. Liu, G. Hua, and T. Chen, “A hierarchical visual model for video object summarization,” *Pattern Analysis and Machine Intelligence*, vol. 32, no. 12, pp. 2178–2190, 2010.
- [5] Y. J. Lee, J. Ghosh, and K. Grauman, “Discovering important people and objects for egocentric video summarization,” in *Computer Vision and Pattern Recognition*, 2012.
- [6] P. Mundur, Y. Rao, and Y. Yesha, “Keyframe-based video summarization using delaunay clustering,” *International Journal on Digital Libraries*, vol. 6, no. 2, pp. 219–232, 2006.
- [7] Y. Hadi, F. Essannouni, and R. O. H. Thami, “Video summarization by k-medoid clustering,” in *ACM Symposium on Applied Computing*, 2006.
- [8] S. E. F. de Avila, A. P. B. a. Lopes, A. da Luz, Jr., and A. de Albuquerque Araújo, “VSUMM: A mechanism designed to produce static video summaries and a novel evaluation method,” *Pattern Recognition Letters*, vol. 32, no. 1, pp. 56–68, 2011.
- [9] M. Furini, F. Geraci, M. Montangero, and M. Pellegrini, “VISTO: Visual storyboard for web video browsing,” in *ACM international conference on Image and video retrieval*, 2007.
- [10] M. Furini, F. Geraci, M. Montangero, and M. Pellegrini, “STIMO: Still and moving video storyboard for the web scenario,” *Multimedia Tools and Applications*, vol. 46, no. 1, pp. 47–69, 2010.
- [11] C.-W. Ngo, Y.-F. Ma, and H.-J. Zhang, “Video summarization and scene detection by graph modeling,” *Circuits and Systems for Video Technology*, vol. 15, no. 2, pp. 296–305, 2005.
- [12] J. Nam and A. H. Tewfik, “Event-driven video abstraction and visualization,” *Multimedia Tools and Applications*, vol. 16, no. 1, pp. 55–77, 2002.
- [13] R. Laganière, R. Bacco, A. Hocevar, P. Lambert, G. Païs, and B. E. Ionescu, “Video summarization from spatio-temporal features,” in *ACM Multimedia*, 2008.
- [14] M. Gygli, H. Grabner, H. Riemenschneider, and L. Van Gool, “Creating summaries from user videos,” in *European Conference on Computer Vision*, 2014.
- [15] Z. Lu and K. Grauman, “Story-driven summarization for egocentric video,” in *Computer Vision and Pattern Recognition*, 2013.
- [16] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” in *Advances in neural information processing systems*, 2012.
- [17] A. S. Razavian, H. Azizpour, J. Sullivan, and S. Carlsson, “CNN features off-the-shelf: An astounding baseline for recognition,” in *Computer Vision and Pattern Recognition*, 2014.
- [18] J. Lin, O. Morère, A. Veillard, V. Chandrasekhar, and H. Goh, “DeepHash: Getting regularization, depth and fine-tuning right,” in *arXiv preprint arXiv:1501.04711*, 2015.
- [19] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell, “Caffe: Convolutional architecture for fast feature embedding,” in *ACM International Conference on Multimedia*, 2014.
- [20] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” in *arXiv preprint arXiv:1409.1556*, 2014.
- [21] B. Zhou, A. Lapedriza, J. Xiao, A. Torralba, and A. Oliva, “Learning deep features for scene recognition using places database,” in *Neural Information Processing Systems*, 2014.
- [22] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “Imagenet: A large-scale hierarchical image database,” in *Computer Vision and Pattern Recognition*, 2009.
- [23] G. E. Hinton, “Training products of experts by minimizing contrastive divergence,” *Neural Computation*, vol. 14, no. 8, pp. 1771–1800, 2002.
- [24] G. E. Hinton, S. Osindero, and Y.-W. Teh, “A fast learning algorithm for deep belief networks,” *Neural Computation*, vol. 18, no. 7, pp. 1527–1554, 2006.
- [25] H. Goh, N. Thome, M. Cord, and J.-H. Lim, “Unsupervised and supervised visual codes with restricted Boltzmann machines,” in *European Conference on Computer Vision*, 2012.
- [26] M. D. Zeiler and R. Fergus, “Visualizing and understanding convolutional networks,” in *European Conference on Computer Vision*, 2014.
- [27] S. Lloyd, “Least squares quantization in PCM,” *Information Theory*, vol. 28, no. 2, pp. 129–137, 1982.